

Porting Code Across Simple Mobile Robots*

Yixiao Lin, Sayan Mitra and Shuting Li
Coordinated Science Laboratory
University of Illinois at Urbana-Champaign

March 3, 2016

Abstract

The StarL programming framework aims to simplify development of distributed robotic applications by providing easy-to-use language constructs for communication and control. It has been used to develop applications such as formation control, distributed tracking, and collaborative search. In this paper, we present a complete redesign of the StarL language and its runtime system which enables us to achieve portability of robot programs across platforms. Thus, the same application program, say, for distributed tracking, can now be compiled and deployed on multiple, heterogeneous robotic platforms. Towards portability, this we first define the semantics of StarL programs in a way that is largely platform independent, except for a few key platform-dependent parameters that capture the worst-case execution and sensing delays and resolution of sensors. Next, we present a design of the StarL runtime system, including a robot controller, that meets the above semantics. The controller consists of a platform-independent path planner implemented using RRTs and a platform-dependent way-point tracker that is implemented using the control commands available for the platform. We demonstrate portability of StarL applications using simulation results for two different robotic platforms, and several applications.

keywords Mobile robots; Programming; Semantics; Runtime system

*.

1 Introduction

Distributed robotic systems have been at the frontier of research in manufacturing [9, 7], transportation [32], logistics [28, 35], and exploration [3]. Yet, managing and experimenting with distributed robotic platforms remains a daunting task. Deploying a new algorithm on a hardware platform typically takes weeks. It takes months to port an existing application program to a new platform with different dynamics or to deploy it on a heterogeneous mix of vehicles. Compare these efforts with the effort needed to perform comparable tasks for desktop applications. This high overhead cost takes a toll on the robustness of the scientific claims and repeatability of experiments.

We are developing the Stabilizing Robotics Language (StarL) [19] and its supporting APIs and runtime system to mitigate some of these issues. StarL provides programming language constructs for coordination and control across robots. See for example, the StarL program for in Section 2.1 for making a collection of robots form a regular polygon. The design goal for this language is to make robot programming faster, more abstract, cleaner, and closer to high-level textbook pseudo-code [21, 8]. The result would be programs that are also easier to debug and verify.

Two key features of StarL are a distributed shared memory (DSM) primitive for coordination and a reach-avoid primitive for control. DSM allows a program to declare program variables that are shared across multiple robots (e.g. `sharedmw x : Int`). This enables programs running on different robots to communicate by writing-to and reading from the shared variable x . The runtime system implementing the DSM ensures that the updated values of x are consistently propagated to appropriate robots using message passing. The reach-avoid primitive provides a function `doReachAvoid(x, U)` and a set of control flags. A StarL program calling `doReachAvoid(x, U)` instructs the low-level robot motion controller to reach the target point x while avoiding the region U (defined in a common coordinate system). The control flags are updated by the low-level controller to indicate status information such as the target has been reached, the region U has been violated, or that the controller is giving-up.

In this paper, we present a newly implemented feature of StarL, namely, portability across different platforms. For some platforms (e.g., an industrial robotic arm vs. a bipedal robot), there is little commonality in the task and their implementations, and therefore, portability is not meaningful. However, there exists a class of distributed robotic platforms and related tasks, such as, visiting a sequence of points in 3D space to achieve some higher-level goal, where automatic portability of can improve both

quality of programs and the productivity of developers. Consider a distributed collaborative search application. The high-level code deals with leader election, allocation of subtasks to robots, and tracking failures and progress. While individual participating robots may have different low-level controllers for tracking way-points, the higher-level code can be made dynamics-independent, and therefore, portable. In this paper, we present the design of the StarL language and supporting runtime system which demonstrate feasibility of this idea. The implementation of our StarL framework is available at <https://github.com/lin187/StarL1.5>

1.1 Overview of contributions

First, we define what it means to port a StarL program across platforms. When two instances $P(m_1)$ and $P(m_2)$ of the same program P are executed on two different desktop computers m_1 and m_2 (with different hardware, OS, and execution environments), we expect the *observable behaviors they produce to be semantically equivalent*, while there may be other aspects of their behaviors (say, execution time) that differ. Analogously, in order to discuss portability of robot programs, first we define the platform-invariant semantics of StarL. Mathematically, we view the evolution of a StarL program together with its hardware stack and physical environment, as a hybrid automaton. Within this framework, we define semantics in terms of the relationship between the calls to *doReachAvoid* and the corresponding changes to the control status flags. A crucial difference in our semantics compared to the semantics of ordinary programs is that, by necessity, our semantics does not completely abstract away timing and position information.

Our second contribution is the design and implementation of a runtime system that enables porting programs across robotic platforms. A StarL program consists of several threads that interact through a part of memory called the variable holder. When a StarL program calls *doReachAvoid*(x, U), it instructs to the robot controller that the new target position is x and a region to be avoided is U . The controller has two continuously executing threads: a path planner and a way-point tracker. The path planner plans possible paths from the current position to the target (avoiding U) and provides a sequence of way-points to the way-point tracker. The path planner may fail to discover a path.

The StarL compiler translates the application program, with additional platform specific information, to executable code for each target platform. The same high-level code could be compiled for different platforms and instrumented for simulations and monitoring. We demonstrate this capability

with several applications and two popular platforms. Our experimental results use the a discrete event simulator that runs actual StarL executable programs (Java) with detailed, physics-based simulations of the platform and the communication channels.

1.2 Related work

Several projects have pushed the limits of custom-built software and hardware for distributed robotics (see [17, 27, 5] for some recent examples); this approach is orthogonal to our objective of developing portable software.

There are several libraries for programming robots using languages like Python and Matlab [26, 12, 2, 6, 24, 25, 4, 31, 29] and notably the Reactive Model-based Programming Language (RMPL) [33], but they do not provide high-level coordination and control APIs nor do they address portability.

The Robotics Operating System (ROS) [26] provides a library of drivers and functions for programming robots and has an ever-growing community of developers and users. While parts of ROS library is reusable, there is currently no effort towards automatically porting programs across platforms. Our previous work on StarL [19, 36] provided software components for mutual exclusion, leader election, service discover, etc. but also did not address portability or the semantic issues.

Our point of view for programming differs from the correct-by-construction synthesis approach (for example, [18, 10, 30, 14, 34, 23, 1]) in what we see as the role of the user. A correct-by-construction synthesis algorithm takes as input a high-level requirement to generate robot programs for accomplishing this requirement. In our approach, human creativity continues to be central in writing the step-by-step program that solves the task, but the tedious and error-prone steps in low-level coordination and control are automated, and specifically in this paper, we focus on enabling automatic porting of programs across platforms.

2 Preliminaries and Problem Setup

2.1 Example: Formation program

The simple formation program in Figure 1 illustrates some of the key features of the StarL language. The program encodes a heuristic which when executed by an (odd) number of robots form a regular polygon. This StarL program is transformed by our compiler to Java program that is executed at

each robot (see, [20] for more details on this transformation). The program has two parts: variable declarations and guarded-commands.

variables When the program is executed by robot i , the scope of the first four variables is only the program at robot i . The variable pid is initialized with a call to the built-in function *getId* which returns a unique identifier for each participant robots. The variable loc is used like a program counter. *ItemPosition* is a built-in type for storing 3D coordinates of points in space with respect to a common and fixed coordinate system.

The **sharedsw** keyword declares $pos[i]$ as a single-writer (sw) multi-reader shared variable. That is, the scope of the integer array $pos[]$ spans all instances of formation program executing, all its components can be read by all instances, but only the instance running at robot i can write to $pos[i]$. The StarL runtime system implements a distributed share memory (DSM) system over wireless message passing. For example, whenever the program at i writes a new value, say 5, to a shared variable x with the assignment $x = 5$, the runtime system detects this change, then sends $\langle x, 5, timestamp \rangle$ to other participants running the same application, and upon receiving these messages, the runtime system at the other end updates its copy of x to 5.

guarded-commands The main part of the program consists of several guarded-command (**pre – eff**) blocks. These blocks are translated to **if** conditionals inside a big **while** loop in the transformed Java program. In each iteration of this big loop, the **effect** part of a block may be executed if its **precondition** holds. If multiple preconditions hold (and no priority order is specified) then one of the effects is executed arbitrarily.

The formation program has three blocks. The *initialize()* block is executed once at the beginning and it sets $pos[i]$ to be the current position of the robot using built-in *getPos* function. It also sets $loc = calc$ which ensures that only the *update()* block can execute next. The next block *update()* is enabled when $loc = calc$ or the special boolean variable *active* is set to **false**. It computes a target position for the robot using the *bisector* function (not shown here). This function when executed by robot i computes a point (at distance len) on the perpendicular bisector of the line joining $pos[j]$ and $pos[(j + 1)\%n]$, where j and $(j + 1)\%n$ are the robots diametrically opposite to i . It is straightforward to check that if $pos[]$ array forms a regular n -sided polygon then the program reaches an equilibrium. The *doReachAvoid(target, 0)* function call communicates to the robot’s controller that the new destination is *target*. Generally, *doReachAvoid(x, U)*

```

1  pid: Int = getId();
   loc: enum{init, calc, wait} = init;
3  counter: Int = 0;
   target: ItemPosition;
5  sharedsw pos[pid]: ItemPosition;

7  initialize()
   pre loc == init
9   eff pos[pid] = getPos(); loc = calc;

11 update()
   pre loc == calc || !active
13  eff target = bisector(pos, pid, len);
   doReachAvoid(target, 0);
15   counter = 0; loc = wait;

17 wait()
   pre loc == wait && counter <= 5
19  eff pos[pid] = getPos();
   counter++;
21  if counter > 5 then loc = calc

```

Figure 1: StarL code of the formation program.

takes two arguments: x is a point in space and U is a set. Roughly, this call communicates to the robot controller to reach x while avoiding U . In the remainder of the paper, we will discuss the precise semantics of *doReachAvoid*, its interaction with StarL applications, and its portable implementations. Finally, the *wait()* block merely updates the $pos[i]$ variable and the *counter* so that the *wait* block is executed 5 times after every execution of *update*.

2.2 Semantics of StarL applications

When a StarL program (such as the one in Figure 1) is ported from one platform to another, we expect that certain important aspects of the behavior of the program to be preserved, while other less crucial aspects may alter. In this section, we discuss the semantics of StarL programs in terms of the expected behavior of certain function calls and variables. This will be the basis for specifying what it means to correctly port StarL programs across different robot platforms.

Semantics of a StarL applications is defined using hybrid I/O automata (HIOA) [22, 16]. HIOA is a mathematical framework for precisely describing systems that have both discrete and continuous dynamics. Discrete dynamics is modeled by *transitions* that instantaneously change the state variables

of the system. Typically transitions are defined by program statements that change state variables associated with the computing stack. Continuous evolution is modeled by *trajectories* that govern how state variables evolve over an interval of time. A trajectory is a function that maps an interval of time to valuations of variables. Typically they are specified using differential and algebraic equations and define the evolution of the physical variables (variables related to the computing stack remain constant over a trajectory).

All the program threads implementing the application, the message channels, as well as the physical environment of the application (robot chassis, obstacles) are modeled as hybrid automata, and the overall system is described by a giant composition of these automata (see Figure 2). The semantics of the control API is then defined in terms of the allowed observable behaviors or traces of **system**. A complete formal presentation of the semantics is beyond the scope of this paper, we refer the interested reader to the publications on virtual nodes which show another application of HIOA framework in giving precise semantics of a distributed programming system [13].

A *hybrid I/O automaton* is a tuple $\mathcal{A} = \langle V, \Theta, A, \mathcal{D}, \mathcal{T} \rangle$, where (i) V is a set of variables partitioned into input, output and state (internal) variables; set of all possible valuations of the state variables Q is the set of states. (ii) $\Theta \subseteq Q$ is a set of initial states, (iii) A is a set of actions partitioned into input, output, and internal actions, (iv) $\mathcal{D} \subseteq Q \times A \times Q$ is the set of discrete transitions, and (v) \mathcal{T} is a well-defined set of trajectories for all the variables V . For any trajectory $\tau \in \mathcal{T}$, $\tau.\text{fstate}$ is the first state and if τ is closed then $\tau.\text{lstate}$ is the last state of τ .

In addition, it is required that \mathcal{A} satisfies the input action and input trajectory enabled condition [22, 16] which implies that it cannot block (a) any input action triggered by another automaton in the system (possibly the environment), and (b) it cannot block any input trajectory produced by another automaton in the system.

A particular behavior of \mathcal{A} is called an *execution* and mathematically it is an alternating sequence $\alpha = \tau_0, a_1, \tau_1, a_2, \dots$, such that $\tau_0.\text{fstate} \in \Theta$ and $(\tau_i.\text{lstate}, a_{i+1}, \tau_{i+1}.\text{fstate}) \in \mathcal{D}$. In general, HIOA are open and nondeterministic, and therefore, have many executions from the same starting state. Once we fix an execution α of an HIOA, we can use the notation $v(t)$ to denote the valuation of a variable $v \in V$ at time t in that execution α .

Each StarL application executes several program threads, and interacts with the environment through different sensors and actuators. Figure 2 shows these different component HIOAs that make-up the model of a StarL application and its environment (the overall HIOA model is obtained by composing the component HIOAs). The HIOA models for the individual

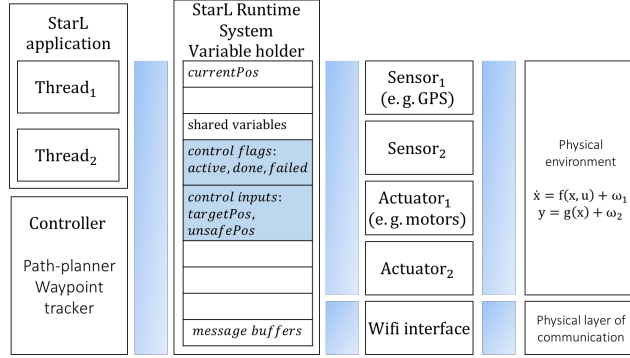


Figure 2: Application threads, controller, variable holder, and physical environment. Shaded areas indicate information flow through messages, sensors, and actuators.

threads capture the evolution of the corresponding program variables according to the guarded commands in the program. Time elapses between the successive transitions, but the exact duration of time is unpredictable and platform dependent. The controller uses special program threads (also modeled as an HIOA) responsible for implementing reach-avoid control strategies. The variable holder is a part of the runtime system that maintains shared variables and message buffers.

All the threads communicate with each other and with the sensors and actuators through variables in the variable holder. For example, the *currentPos* variable is updated periodically by the positioning sensors (e.g. GPS) and read by the application programs by calling the *getPos()* function. For this paper, we will focus on the control API of StarL which consists of the following single-writer multi-reader shared variables stored in the variable holder:

Each of these variables are written to by a single thread as indicated in the table. When a call to *doReachAvoid(x,U)* from the application program is executed, it sets:

$targetPos := x;$
 $unsafePos := U;$

thus communicating to the controller thread that the new target is x and the unsafe region is U . Once this call is executed, the control thread attempts to move the robot towards x while avoiding U and updates the three status variables *active*, *done*, and *failed*.

With this set-up, we can now define the key semantic properties of StarL

Shared variable	Writer	Interpretation
<i>targetPos</i>	App.	coordinates of target position
<i>currentPos</i>	Sensor	coordinates of last recorded position
<i>unsafePos</i>	App.	set of coordinates of unsafe positions
<i>active</i>	Ctrl.	attempting progress towards target
<i>done</i>	Ctrl.	reached vicinity of target
<i>failed</i>	Ctrl.	cannot make progress or entered unsafe

Table 1: Single-writer multi-reader shared variables in the control API. Each of the variables are written to by one of the following threads: the main application thread (App.), the controller thread (Ctrl.), and sensor update threads.

programs. In all of the following statements, we fix any execution α of the HIOA describing the system, we assume that at some time $t_0 > 0$ a *doReachAvoid*(x, U) call completes execution and there are no other calls to *doReachAvoid* between t_0 and another arbitrary time point $T > t_0$. Semantics of the control flags are specified in terms of two platform dependent parameters: a dwell time $d_t > 0$ and a quantization distance $q_d > 0$. We define the predicate $reached(t) := d(currentPos(t_1), targetPos(t)) \leq q_d$, that is, the current recorded position of the robot in the variable holder is within q_d of the last target position issued by invocation of *doReachAvoid*.

$$\begin{aligned}
D1 : & \exists t \in [t_0, T] : done(t) \implies \exists t_1 \in [t_0, t] : reach(t_1). \\
D2 : & \exists t_1 > t_0, \forall t \in [t_1, t_1 + d_t], reach(t) \\
& \implies \exists t_2 > t_1, \forall t' \in [t_2, T], done(t').
\end{aligned}$$

Note the asymmetry in the two conditions. D1 states that along any execution *reach* precedes the corresponding *done*. D2 states that along any execution, if *reach* persists for at least d_t duration (and no further *doReachAvoid*'s are issued) then *done* becomes set. Persistence is required for the runtime system to detect predicates. Next, we define the predicate $crossed(t) := d(currentPos(t_1), unsafe(t)) \leq q_d$, that is, the current recorded position of the robot is within q_d distance of an unsafe position issued by invocation of *doReachAvoid*.

$$\begin{aligned}
F1 : & \exists t \in [t_0, T], failed(t) \implies \exists t_1 \in [t_0, t], crossed(t) \\
F2 : & \exists t_1 > t_0, \forall t \in [t_1, t_1 + d_t], crossed(t) \\
& \implies \exists t_2 > t_1, \forall t \in [t_2, T], failed(t).
\end{aligned}$$

F1 states that along any execution *crossed* precedes the corresponding *failed*. F2 states that along any execution, if *crossed* persists for at least d_t

duration (and no further *doReachAvoid*'s are issued) then *failed* becomes set. In addition to the *done* and *failed* flags, the controller API provides the *active* status flag for the low-level controllers to communicate with the application programs.

$$A1 : \exists t \in [t_0, T], active(t) \implies \forall t_1 \leq t, \neg done(t_1) \wedge \neg failed(t_1).$$

A1 implies that at any point in an execution, if *active* is set then it is neither preceded by *done* nor by *failed*. The semantics of *active* being set to false is unspecified. Informally, this means that the low-level controller has given-up on attempting to reach the *targetPos*, but the precise conditions under which this can happen may depend on the platform, the implementation of the controllers, and the environment which may include moving obstacles.

2.3 Semantics of porting

When a StarL application is ported correctly from one platform to another, we expect the above conditions to be preserved albeit with different (known) values of the platform dependent constants d_t and q_d . We have decided to state the semantics in terms of program statements and the control flags. This leaves a necessary gap between these statements and the actual physical state of the robot. (1) The *done* and *failed* flags imply conditions only on the recorded position of the robot and not its physical position. Under additional assumptions about the veracity of sensors and underlying dynamics, we can infer that indeed *done* implies that the actual position of the robot is q_d -close to the target issued. (2) Conversely, under additional assumptions about persistence of the actual position of the robot near the target and the speed of the runtime system, we can infer that if the real position is close to the target then *done* will be set. Similar arguments can be constructed for reasoning about the *unsafe* set and the *failed* flag.

In the next section, we describe our design of the StarL controller thread and runtime system which enables us to correctly and automatically port StarL applications across different platforms.

3 Implementation

A controller for a robot gets inputs from other StarL programs through the variables *targetPos* and *unsafePos*, and produces outputs by writing to the status variables *done*, *failed* and *active*. It may receive input data from

other variables and sensors as well. A valid controller should guarantee the semantics given in the previous section. In this section, we describe our implementation choices. Our implementation involves two separate threads: (a) a path planner and (b) a way-point tracker.

3.1 Way-point tracker

The way-point tracker reads inputs *currentPos* and a way-point (generated by the path planner), and uses low-level control instructions and a dynamical model of the robot, to attempt closed-loop tracking of the way-point. It is oblivious to obstacles and the unsafe set. Like device drivers, this program has to be written for each platform independently as both the low-level control instructions and the dynamics of the model are strongly platform dependent. For example, the low-level control instructions and the dynamical models for the AR Drone2 and the iRobot Create are shown in Tables 2 and 3. The commands for setting YawSpeed, Pitch, Roll, Gaz can be used in any combination while other commands can only be used one at a time. For the AR Drone2 model, the state variables are x, y, z position coordi-

AR Drone2	iRobot Create
takeOff()	straight(v_{ref})
land()	turn(a_{ref})
hover()	curve(v_{ref}, r)
setYawSpeed(a_{ref})	
setPitch(θ_{ref})	
setRoll(ϕ_{ref})	
setGaz(v_{ref})	

Table 2: Low-level control instructions available for ARDrone2 and the iRobot Create.

nates, yaw (ψ), pitch (θ), roll (ϕ) angles and the corresponding velocities. For the iRobot model the state variables are x, y position coordinates and the heading (θ) angle.

AR Drone2	iRobot Create
$thrust = (gaz + 10) / \cos \phi / \cos \theta$	$\dot{x} = v \cdot \cos \theta$
$\ddot{x} = -(thrust)(\sin \phi \cdot \sin \psi + \cos \phi \cdot \sin \theta \cdot \cos \psi) / M$	$\dot{y} = v \cdot \sin \theta$
$\ddot{y} = (thrust)(\sin \phi \cdot \cos \psi + \cos \phi \cdot \sin \theta \cdot \sin \psi) / M$	$\dot{\theta} = a_{\text{ref}}$
$\dot{z} = gaz$	
$\dot{\psi} = gain \cdot (a_{\text{ref}} - \psi)$	

Table 3: Dynamical models of AR Drone2 and iRobot Create.

Since control commands are provided by the robotic platforms as black boxes, in creating the model for these controller we use standard proportional controllers. For example, the yaw speed control is shown in the last

row and similar controls for pitch and roll are used (omitted in the table). This technique can be extended to other closed loop dynamics such as PID. The way-point tracker reads sensor inputs and uses the motion commands in a closed-loop.

3.2 Safe path planner

The path planner attempts to generate a sequence of way-points from *currentPos* leading to *targetPos* such that if the robot follows this path it will avoid the set of locations specified by *unsafePos*. It uses a platform-independent implementation of RRT algorithm [15], instantiated with platform dependent way-point controller.

Whenever a change in the *targetPos* or *currentPos* is detected, the following algorithm for growing an RRT is executed with an initial tree that has a single node at *currentPos*. Given *currentPos*, *targetPos*, *unsafePos* and a current tree T in \mathbb{R}^3 , the RRT algorithm adds new points to the tree as follows: It picks a random point $\vec{x} \in \mathbb{R}^3$ and attempts to add it to T . In doing so, it first picks a point in the tree $\vec{p} \in T$ that is closest to \vec{x} . Then it simulates the way-point tracker with a model of the plant to go from \vec{p} to \vec{x} to generate a simulated path. Finally, it checks if this simulated path is sufficiently far from *unsafePos*. If this check succeeds then it adds \vec{x} to the tree. Otherwise, it finds another point halfway between \vec{p} and \vec{x} unless the distance between the two falls below a threshold in which case a new \vec{x} is picked at random.

In this way, once the tree T reaches a threshold size the RRT construction algorithm is stopped. If a path from the vicinity of *currentPos* to the vicinity of *targetPos* exists in T then this path is sent to the way-point tracker. Otherwise, the *active* flag is set to *false*. Assuming that *currentPos* and *targetPos* did not change during the tree construction, this indicates to the StarL application program that a safe path has not been discovered by the path planner. In general, establishing nonexistence of safe path is challenging, and our design of the runtime system leaves it to the programmer to code best effort strategies by detecting when all the control flags *active*, *done*, and *failed* simultaneously become *false*.

4 Simulation experiments

StarL applications have been deployed on two different types of robotic platforms at UIUC and UT Arlington: iRobot Create ground robots controlled

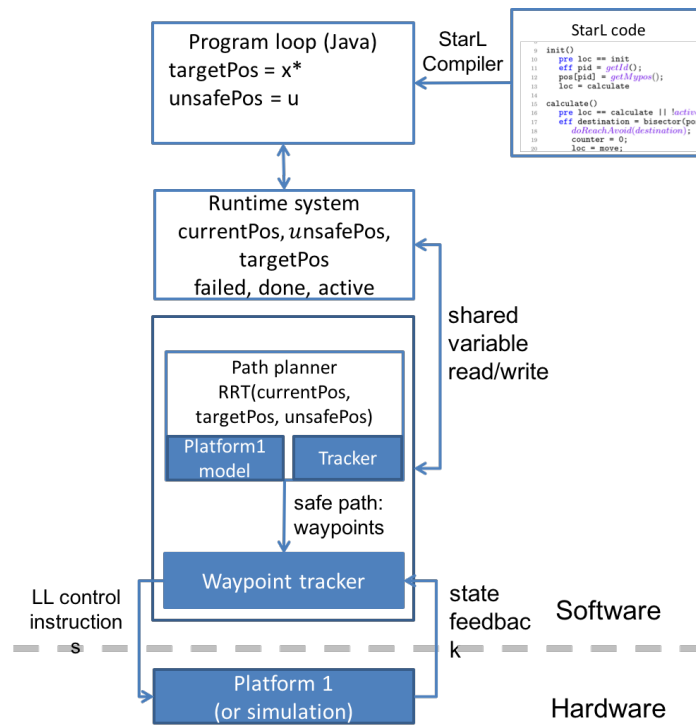


Figure 3: Porting StarL applications. The shaded blocks are platform dependent and the rest are generic.

by Android smart phones [11] and AR Drone quadrotors¹. StarL also has a powerful discrete event simulator that can run the same application code with dozens of robot instances, a simulated physical environment with detailed physics models, and with different kinds of robots. In this section, we report on semantics and portability of StarL applications with simulation experiments.

To port a StarL application to a new platform, developers need to have the appropriate drivers for the hardware (e.g. wifi interfaces, sensors, positioning system, etc.), write the way-point tracker, and provide the physics model of the platform. These are the platform-dependent components shown in blue in Figure 3. For our iRobot+Android platform, for example, the smartphones handle the computation and communication (using UDP sockets) and read the sensor data from the iRobot sensors, and write actuation commands to the motion controllers.

4.1 Experiments on semantics of ReachAvoid

Our first set of experiments illustrate the semantics of *doReachAvoid* for individual robots in a static environment. A single robot executes a simple program that is a sequence of four *doReachAvoid*($x[i], U$) calls. That is, once the *done* flag is set after the first call, then *doReachAvoid*($x[i + 1], U$) is called and so on. Here U is a wall as shown in Figure 4 and we call the four points A, B, C , and D .

Two independent traces of running this StarL program on the iRobot (red) and the AR Drone (blue) models are shown in Figure 4. The hollow circles labeled $Rit = xx.xs$ indicate that the time when the i^{th} *doReachAvoid*($x[i], U$) was executed by the robot. As expected, the AR Drone is slow to get to A as it has to first take-off, but then it reaches B and C much sooner.

Figure 5 shows the time versus x-position plots for four different runs of the same program on two different platforms. There are several sources of uncertainty that make every simulation trace different, even on the same platform. For example, the RRT algorithm used by the path-planner is randomized, there is uncertainty in the scheduling of different threads, the underlying simulation models for the dynamics involves random effects such as wind, sensor noise, etc. Across different platforms the timing parameter d_t and the quantization distance q_d are different, and of course, the dynamics is also different.

For each of these traces, once again, the hollow circles show the time

¹<https://www.youtube.com/watch?v=27NCnLYHAtY>




Figure 4: Two independent traces of a quadcopter (blue) and a iRobot (red) following a fixed sequence of four way-points (solid circles) with a fixed unsafe set (rectangle). The hollow circles are points where *doReachAvoid* is called.

of the *doReachAvoid*($x[i], U$) call and the dashed line to its right show the x -coordinate of $x[i]$, that is, the next *targetPos*. Once the robot reaches the q_d -ball around this target position for both x and the y coordinate (latter is not shown in the Figures), then the *done* flag is set and soon afterwards (depending on the execution time of the program thread) a new *doReachAvoid*($x[i], U$) is called.

Observe that the quadcopter overshoots the *targetPos* sometimes (B in Figure 4). The black execution of the quadcopter ends with the *fail* flag being set as the quadcopter drifts into *unsafePos* because of wind (Figure 5).

If there are moving objects blocking progress or if the *targetPos* is too close to the *unsafePos*, the path planner may be unable to find a safe path in which case it sets the *active* flag to false to inform the application program.

4.2 Discussion and other applications

While a detailed user study has been beyond the scope of the current work, we believe that the StarL language constructs, primitives, and now the run time system supporting portability, dramatically improves programmer productivity and experience for developing distributed mobile robotic applications.

To illustrate this, we discuss some other StarL applications. Figure 6 shows the code for a distributed way-point tracking application in which the participating robots need to collectively cover a stream of way-points.




Figure 5: Timed traces of quadcopters (top) and iRobot (bottom). The x-coordinate (solid lines), calls to `doReachAvoid` (circles), and the target x-coordinates (dashed lines). One of the quadcopter executions end with the *failed* flag being set.

The shared variable *sharedIndex* is used by the group to track the current way-point of interest. Each robot reads the *sharedIndex* and stores it in the local variable *currentIndex*. Then it calls the *doReachAvoid* function to race to the target while avoiding obstacles. If and when the *done* flag becomes true, *sharedIndex* is incremented which indicates to all the other robots that its time to move to the next way-point. If the shared variable is changed, the program again calls *doReachAvoid* function to move to the next target.

In addition to reach-avoid and the DSM, StarL also provides primitives that are commonly used idioms in distributed systems such as leader election, mutual exclusion, and set consensus. Programs for complex applications become simple and modular using these primitives. A number of portable demo applications for distributed robots have been developed [19]. For example, our collaborative search builds up on the distributed way-point


```

    loc: enum{init, pick, wait} = init;
2 target[] : ItemPositon
    currentIndex: int;
4 sharedmw sharedIndex: int = 0;

6 initialize()
    pre loc == init
8    eff target = getInput(); currentIndex = 0; loc = pick;

10 pick()
    pre sharedIndex < targets.size() &&
12    (loc == pick || !active || failed || currentIndex != sharedIndex )
    eff currentIndex = sharedIndex;
14    doReachAvoid(target[currentIndex], getObstacles());
        loc = wait;
16
18 wait()
    pre loc == wait && done
    eff sharedIndex = currentIndex + 1; loc = pick;

```

Figure 6: Distributed way-point tracking application.

tracking applications: The input is a map with a number of rooms to be searched. At the beginning, a leader is elected; the leader assigns the rooms (way-points) to all the robots using DSM. When each robot receives the assignment from the leader, it uses *doReachAvoid* to move to the entrance of the room and then search the room. Finally, it updates the DSM to relay the status of the search to the others. A screen shot of the StarL simulator running this application is shown in Figure 7.

5 Conclusions

In this paper, we introduce the new StarL programming framework which supports porting of high-level application programs across different platforms. We showed utility of the language constructs for control and coordination, namely *doReachAvoid* and distributed shared memory, through several example applications. In order to discuss portability of StarL programs, we first propose a semantics of these programs. It is meaningless to discuss semantics of a robotic program without specific reference to timing and spatial positions and orientations. At the same time, if the semantics is too closely attached to the specific dynamics of a platform, the purpose of portability is defeated. In this paper, we reach a compromise and propose the semantics of StarL applications in terms of key program variables




Figure 7: Screen shot of Distributed Search Simulation

that are platform independent, and certain platform dependent parameters related to sensing precision and timing. This semantics then enables us to state what it means for an application program to be ported to a new platform with new parameters but identical program interfaces. Based on this semantics, we present the design of the StarL runtime system that enables us to automatically port application programs. Specifically, given the dynamical model of a target platform and a low-level controller interface, the StarL compiler automatically generates executable code for the platform that meets the semantics. The key element of the runtime system is a robot controller that interfaces with the `doReachAvoid` construct. Through detailed simulations of StarL application programs on two popular platforms, we have established that the runtime system can meet the proposed semantics.

6 Acknowledgments

We are thankful to individuals who have contributed to the development and evaluation of the StarL platform. Ritwika Ghosh developed the compiler for the high-level language. Zhenqi Huang and Kenji Fukuda developed the way-point tracker for the AR Drone2 platform.

References

- [1] Tewodros A. Beyene, Swarat Chaudhuri, Corneliu Popeea, and Andrey Rybalchenko. A constraint-based approach to solving games on infinite graphs. In *The 41st Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL '14, San Diego, CA, USA, January 20-21, 2014*, pages 221–234, 2014.
- [2] Douglas Blank, Lisa Meeden, and Deepak Kumar. Python robotics: An environment for exploring robotics beyond legos. In *ACM SIGCSE Bulletin*, volume 35, pages 317–321. ACM, 2003.
- [3] Jennifer L Burke, Robin R Murphy, Michael D Coover, and Dawn L Riddle. Moonlight in miami: Field study of human-robot interaction in the context of an urban search and rescue disaster response training exercise. *Human-Computer Interaction*, 19(1-2):85–116, 2004.
- [4] Daniele Calisi, Andrea Censi, Luca Iocchi, and Daniele Nardi. Openrdk: a modular framework for robotic software development. In *Intelligent Robots and Systems, 2008. IROS 2008. IEEE/RSJ International Conference on*, pages 1872–1877. IEEE, 2008.
- [5] Luiz Chaimowicz, Nathan Michael, and Vijay Kumar. Controlling swarms of robots using interpolated implicit functions. In *Robotics and Automation, 2005. ICRA 2005. Proceedings of the 2005 IEEE International Conference on*, pages 2487–2492. IEEE, 2005.
- [6] Peter Corke. A robotics toolbox for matlab. *Robotics & Automation Magazine, IEEE*, 3(1):24–32, 1996.
- [7] Nikolaus Correll. Parameter estimation and optimal control of swarm-robotic systems: A case study in distributed task allocation. In *2008 IEEE International Conference on Robotics and Automation, ICRA 2008, May 19-23, 2008, Pasadena, California, USA*, pages 3302–3307, 2008.
- [8] J. Cortes, S. Martinez, T. Karatas, and F. Bullo. Coverage control for mobile sensing networks. In *Robotics and Automation, 2002. Proceedings. ICRA '02. IEEE International Conference on*, volume 2, pages 1327–1332, 2002.
- [9] David M Diltz, Neil P Boyd, and HH Whorms. The evolution of control architectures for automated manufacturing systems. *Journal of manufacturing systems*, 10(1):79–93, 1991.

- [10] Jerry Ding, Eugene Li, Haomiao Huang, and Claire J. Tomlin. Reachability-based synthesis of feedback policies for motion planning under bounded disturbances. In *IEEE International Conference on Robotics and Automation, ICRA 2011, Shanghai, China, 9-13 May 2011*, pages 2160–2165, 2011.
- [11] Parasara Sridhar Duggirala, Taylor T. Johnson, Adam Zimmerman, and Sayan Mitra. Static and dynamic analysis of timed distributed traces. In *Proceedings of the 33rd IEEE Real-Time Systems Symposium, RTSS*, pages 173–182, 2012.
- [12] Brian Gerkey, Richard T Vaughan, and Andrew Howard. The player/stage project: Tools for multi-robot and distributed sensor systems. In *Proceedings of the 11th international conference on advanced robotics*, volume 1, pages 317–323, 2003.
- [13] Seth Gilbert, Nancy Lynch, Sayan Mitra, and Tina Nolte. Self-stabilizing robot formations over unreliable networks. *Special Issue on Self-adaptive and Self-organizing Wireless Networking Systems of ACM Transactions on Autonomous and Adaptive Systems (TAAS)*, 4(3):1–29, 2009.
- [14] Ebru Aydin Gol, Xu Chu Ding, Mircea Lazar, and Calin Belta. Finite bisimulations for switched linear systems. *IEEE Trans. Automat. Contr.*, 59(12):3122–3134, 2014.
- [15] S. Karaman and E. Frazzoli. Sampling-based optimal motion planning for non-holonomic dynamical systems. In *Robotics and Automation (ICRA), 2013 IEEE International Conference on*, pages 5041–5047, May 2013.
- [16] Dilsun K. Kaynar, Nancy Lynch, Roberto Segala, and Frits Vaandrager. *The Theory of Timed I/O Automata*. Synthesis Lectures on Computer Science. Morgan Claypool, November 2005. Also available as Technical Report MIT-LCS-TR-917.
- [17] Kurt Konolige, Charlie Ortiz, Regis Vincent, Benoit Morisset, Andrew Agno, Michael Eriksen, Dieter Fox, Benson Limketkai, Jonathan Ko, Benjamin Stewart, et al. Centibots: Very large scale distributed robotic teams. In *Building the Information Society*, pages 761–761. Springer, 2004.

- [18] Hadas Kress-Gazit, Georgios E. Fainekos, and George J. Pappas. Temporal-logic-based reactive mission and motion planning. *IEEE Transactions on Robotics*, 25(6):1370–1381, 2009.
- [19] Yixiao Lin and Sayan Mitra. Starl: Towards a unified framework for programming, simulating and verifying distributed robotic systems. In *Proceedings of the 16th ACM SIGPLAN/SIGBED Conference on Languages, Compilers and Tools for Embedded Systems, LCTES 2015, CD-ROM, Portland, OR, USA, June 18 - 19, 2015*, pages 9:1–9:10, 2015.
- [20] Yixiao Lin and Sayan Mitra. Starl: Towards a unified framework for programming, simulating and verifying distributed robotic systems. *CoRR*, abs/1502.06286, 2015.
- [21] M. Mesbahi and Magnus Egerstedt. *Graph-theoretic Methods in Multi-agent Networks*. Princeton University Press.
- [22] Sayan Mitra. *A Verification Framework for Hybrid Systems*. PhD thesis, Massachusetts Institute of Technology, Cambridge, MA 02139, September 2007.
- [23] Srinivas Nedunuri, Sailesh Prabhu, Mark Moll, Swarat Chaudhuri, and Lydia E. Kavraki. Smt-based synthesis of integrated task and motion plans from plan outlines. In *2014 IEEE International Conference on Robotics and Automation, ICRA 2014, Hong Kong, China, May 31 - June 7, 2014*, pages 655–662, 2014.
- [24] IA Nesnas. Claraty: A collaborative software for advancing robotic technologies. In *Proc. of NASA Science and Technology Conference*, volume 2, 2007.
- [25] Issa AD Nesnas. The claraty project: Coping with hardware and software heterogeneity. In *Software Engineering for Experimental Robotics*, pages 31–70. Springer, 2007.
- [26] Morgan Quigley, Ken Conley, Brian Gerkey, Josh Faust, Tully Foote, Jeremy Leibs, Rob Wheeler, and Andrew Y Ng. Ros: an open-source Robot Operating System. In *ICRA workshop on open source software*, volume 3, page 5, 2009.
- [27] Michael Rubenstein, Christian Ahler, and Radhika Nagpal. Kilobot: A low cost scalable robot system for collective behaviors. In *Robotics and Automation (ICRA), 2012 IEEE International Conference on*, pages 3293–3298. IEEE, 2012.

- [28] Christopher Steiner. Bot in the delivery:kiva systems. *Forbes Magazine*, March 2009. http://www.forbes.com/forbes/2009/0316/040_bot_time_saves_nine.html.
- [29] Andrew Stubbs, Vladimeros Vladimerou, Adam Thomas Fulford, Derek King, Jeffrey Strick, and Geir E Dullerud. Multivehicle systems control over networks: a hovercraft testbed for networked and decentralized control. *Control Systems, IEEE*, 26(3):56–69, 2006.
- [30] Maria Svorenova, Jan Kretínský, Martin Chmelik, Krishnendu Chatterjee, Ivana Cerná, and Calin Belta. Temporal logic control for stochastic linear systems using abstraction refinement of probabilistic games. *CoRR*, abs/1410.5387, 2014.
- [31] Vladimeros Vladimerou, Andrew Stubbs, Joel Rubel, Adam Fulford, and GE Dullerud. A hovercraft testbed for decentralized and cooperative control. In *Proc. of American Control Conference*, 2004.
- [32] M. Mitchell Waldrop. Autonomous vehicles: No drivers required, February 2015.
- [33] Brian C Williams, Michel D Ingham, Seung H Chung, and Paul H Elliott. Model-based programming of intelligent embedded systems and robotic space explorers. *Proceedings of the IEEE*, 91(1):212–237, 2003.
- [34] Tichakorn Wongpiromsarn, Ufuk Topcu, and Richard M. Murray. Receding horizon temporal logic planning. *IEEE Trans. Automat. Contr.*, 57(11):2817–2830, 2012.
- [35] Peter R. Wurman, Raffaello D’Andrea, and Mick Mountz. Coordinating hundreds of cooperative, autonomous vehicles in warehouses. *AI Magazine*, 29(1):9–20, 2008.
- [36] Adam Zimmerman. Starl for programming reliable robotic networks. Master’s thesis, Department of Electrical and Computer Engineering, University of Illinois at Urbana Champaign, 2013.